ABRAMS'S STABLE EQUIVALENCE FOR GRAPH BRAID GROUPS

PAUL PRUE AND TRAVIS SCRIMSHAW

Abstract. In his PhD thesis [1], Abrams proved that for a natural number n and a graph G with at least n vertices, the n-strand configuration space of G deformation retracts to a compact subspace, the *discretized* n-strand configuration space, provided G satisfies the following two conditions. First, each path between distinct *essential vertices* (vertices of valence not equal to 2) is of length at least $n+1$ edges, and second, each path from a vertex to itself which is not nullhomotopic is of length at least $n+1$ edges. We prove the first condition can be relaxed to require only that each path between distinct essential vertices is of length at least $n-1$. In doing so, we fill a minor hole in Abrams's original proof. We show the improved result is optimal; that is, the conditions are in fact necessary for the existence of the indicated deformation retraction.

1. Introduction

The goal of this paper is to characterize exactly when a *braid group on a graph* may be studied via a certain CW complex associated to the graph. Let X denote a connected topological space. An *n-strand configuration in X* is an n-element subset of X. The *unordered n-strand configuration space* of X is the space of unordered subsets consisting of n distinct elements of X, and is denoted $U\mathcal{C}^n(X)$. (We use the terms *labeled* and *unlabeled* as synonyms for the terms *ordered* and *unordered*, respectively.) For a positive integer n, the classical braid group B_n is the fundamental group $\pi_1(U\mathcal{C}^n(D^2))$, where D^2 is the 2-dimensional topological disk. Thus, from the configuration-space perspective, a braid is simply a loop in the space $U\mathcal{C}^n(D^2)$. Similarly, the *ordered n-strand configuration space* of D^2, denoted $\mathcal{C}^n(D^2)$, is the space of ordered tuples consisting of n distinct elements of X. The classical *n-strand pure braid group*, denoted PB_n, is the fundamental group of the ordered

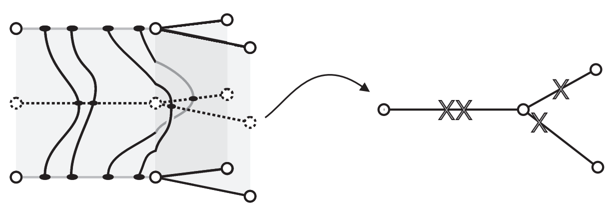

Figure 1: A nontrivial 4-braid in the cylinder $K_{3,1} \times [0,1]$. At each $t \in [0,1]$, an n-braid defines a configuration of n points of the graph, illustrated by the X's on the graph at right.

n-strand configuration space of a disk. Note, the quotient map from the configuration space $\mathcal{C}^n(D^2)$ to the unordered configuration space $U\mathcal{C}^n(D^2)$ induces a short exact sequence $1 \to PB_n \to B_n \to \Sigma_n \to 1$, where Σ_n is the symmetric group on n symbols. For an extensive reference on classical braid groups, see [3].

In the case of graph braid groups, we let $X = G$ be a connected graph, viewed as a 1-dimensional CW complex. The ordered n-strand configuration space of G is denoted $\mathcal{C}^n(G)$. The *n-strand pure graph braid group $PB_n(G)$* is the fundamental group $\pi_1(\mathcal{C}^n(G))$. The unordered configuration space is $U\mathcal{C}^n(G)$, and its fundamental group $\pi_1(U\mathcal{C}^n(G))$ is the *n-strand graph braid group $B_n(G)$*. Note that these fundamental groups do not depend on basepoint: usually the configuration space is connected, but even when it is disconnected the components are all homeomorphic [1].

Graph braid groups, like classical braid groups, can also be viewed as isotopy classes rel endpoints of *braids* (i.e., certain n-tuples of pairwise-disjoint paths) in the cylinder on a topological space. For classical braid groups, this cylinder is $D^2 \times I$, where I is the interval $[0, 1]$. In the case of graph braid groups, one considers instead braids in the cylinder $G \times I$. Figure 1 shows a non-trivial braid in $G \times I$, where G is isomorphic to the complete bipartite graph $K_{3,1}$. A braid $\beta : I \to G \times I$ can be thought of as describing the simultaneous and continuous movements of the n strands, or *tokens*, without collisions, on G. At each $t \in I$, the expression $\beta(t)$ defines a configuration of the n tokens on G. Since β is a loop in the (ordered / unordered) configuration space, it follows that the configurations $\beta(0)$ and $\beta(1)$ are equal.

Besides providing a class of interesting mathematical objects, graph braid groups have real-world applications that have been discussed in [2] and [9]. An example often given is that of a fleet of mobile robots inside a factory, whose movement is confined to a network of track or guide tape. For an idealized robot of infinitesimal size, the configuration space of points on a graph shaped like the track network exactly describes the space of configurations of the fleet of robots.

Abrams introduces the notion of a *discretized configuration space* of n strands on a graph G in his PhD thesis [1]. This is a compact subspace of a configuration space of G, consisting of only those n-strand configurations x in which, for each pair of strands in x, every path in G between the two strands contains at least one full edge of G. Note that $\mathcal{C}^n(G)$ is a subspace of the cubical complex $\prod^n G$, but does not inherit its CW structure, as it is not a compact subspace. In contrast, the *discretized labeled configuration space* of a graph G, denoted $\mathcal{D}^n(G)$, has a CW complex structure as a subcomplex of $\prod^n G$, as does the *discretized unlabeled configuration space*, denoted $U\mathcal{D}^n(G)$. The latter space is obtained as a quotient of the former by the action of the symmetric group Σ_n, which acts by permuting the coordinates of a configuration. Interesting examples among these discretized configuration spaces have been described by Abrams and Ghrist (see [1], [2], and [9]), and we will consider additional examples in this paper.

For a given n, Theorem 2.1 of [1] gives sufficient conditions on G to guarantee that $\mathcal{C}^n(G)$ deformation retracts onto the subspace $\mathcal{D}^n(G)$. We

state it here for reference. An *essential vertex* of a graph is a vertex whose degree is not equal to 2.

Theorem A. [1] Let G be a graph with at least n vertices. Then $\mathcal{C}^n(G)$ deformation retracts onto $\mathcal{D}^n(G)$ if

(A′) each path connecting distinct essential vertices of G has length at least $n+1$, and
(B′) each homotopically essential path connecting a vertex to itself has length at least $n+1$.

The homotopy equivalence established by Theorem A guarantees that one obtains the n-strand pure braid group of G as $\pi_1(\mathcal{D}^n(G))$, provided G is subdivided as in the hypotheses of the theorem. However, condition (A′) of Abrams's theorem is not best possible, as he observes. Several papers (including [2], [7], and [10]) have cited the theorem incorrectly, assuming the optimal bounds on subdivision for which we give proof.

In this paper, we prove Abrams's conjectured optimal bounds on graph subdivision such that the n-strand configuration space of a graph G deformation retracts to the n-strand discretized configuration space. This is the Stable Equivalence Theorem stated here. We write $|G|$ to denote the number of vertices in G.

Theorem 1 (Stable Equivalence). Let $n > 1$ be an integer. For a connected graph G with $|G| > 1$, the space $\mathcal{C}^n(G)$ deformation retracts onto $\mathcal{D}^n(G)$ *if and only if*

(A) each path connecting distinct essential vertices of G has length at least $n-1$, and
(B) each homotopically essential path connecting a vertex to itself has length at least $n+1$.

The result in the theorem applies in the unlabeled case as well, precisely characterizing those pairs n and G for which $U\mathcal{C}^n(G)$ deformation retracts onto $U\mathcal{D}^n(G)$. Note that conditions (A) and (B) are sufficient (and necessary) for the existence of a deformation retraction from the configuration space to the discretized space. Thus, the following definition of *sufficient subdivision* in G for a specified n is optimal.

Definition 1.1. Let G be a finite, connected graph, and let the integer $n > 1$ be given. We say G is *sufficiently subdivided for n tokens* if and only if G satisfies conditions (A) and (B) of Theorem 1.

Many results about graph braid groups are known. In [9], Ghrist showed that the unordered configuration spaces $U\mathcal{C}^n(G)$ are Eilenberg-Maclane spaces of type $K(B_n(G), 1)$. Abrams showed in [1] that graph braid groups have solvable word and conjugacy problems, as they are fundamental groups of locally CAT(0) cubical complexes. Farley and Sabalka use Forman's discrete Morse theory in [7] and [8] to calculate presentations for graph braid groups. In [5], Farley showed that braid groups of trees have free abelian integral homology groups in every dimension. In [6], Farley computes presentations of the integral cohomology rings $H^*(U\mathcal{C}^n(G); \mathbb{Z})$, where T is any tree and n is arbitrary. Crisp and Wiest showed [4] that every graph braid group

embeds in a right-angled Artin group, hence, graph braid groups are linear, bi-orderable, and residually finite.

Our proof of Theorem 1 closely follows the method used in the proof of Theorem A in [1]. The proof relies on two lemmas to show that when a graph G contains the specified subdivision, the inclusion $i : \mathcal{D}^n(G) \to \mathcal{C}^n(G)$ induces a bijection on the homotopy groups π_0 and π_1. We shall sometimes refer to these as the Path Homotopy Lemma (2.1) and the Contractible Loop Lemma (2.2). These are Lemmas 2.2 and 2.3 in [1]. The proof given there for the Contractible Loop Lemma contains a slight error, which we correct on page 16. In our proof of Lemma 2.2, we keep very careful track of the additional standardized configurations and paths in $\mathcal{D}^n(G)$ necessary to account for the relaxed subdivision in the hypotheses of Theorem 1. To prove the other direction in the theorem, we formalize the proof sketch given in [1] to demonstrate that the graph subdivision described in Theorem 1 is in fact necessary for the existence of a deformation retraction.

The outline of this paper is as follows. In Section 2, we prove Theorem 1. Section 3 contains two examples of unlabeled discretized configuration spaces of graphs optimally subdivided for three strands. In both examples, we illustrate the space and compute the graph braid group. In Section 4, we prove the Path Homotopy Lemma and the Contractible Path Lemma.

The authors gratefully acknowledge Lucas Sabalka for his mentorship and encouragement. This paper grew out of a summer 2008 REU with him, supported by an NSF VIGRE grant at UC Davis.

2. Proof of Theorem 1

The proof of Theorem 1 uses the Path Homotopy Lemma and the Contractible Path Lemma, so we state these now.

Lemma 2.1 (Path Homotopy Lemma). Let $\alpha : (I, \partial I) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$. If conditions (A) and (B) of Theorem 1 hold then α can be homotoped rel ∂I into $\mathcal{D}^n(G)$.

Lemma 2.2 (Contractible Loop Lemma). Let $f : (D^2, \partial D^2) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$. If conditions (A) and (B) of Theorem 1 hold then there exists $g : D^2 \to \mathcal{D}^n(G)$ such that $g|_{\partial D^2} = f|_{\partial D^2}$.

We defer the proofs of the two lemmas to Section 4.

The homotopy equivalence established by Theorem 1 ensures the graph (pure) braid group is still obtained as a fundamental group after the discretization, provided the graph satisfies the given conditions. For two examples of this process, see Section 3.

Proof of Theorem 1.

($\Leftarrow$)

Since $n > 1$, G contains at least one essential path (or at least one cycle). By hypothesis, that path (cycle) contains at least n vertices. As mentioned in the Introduction, $\mathcal{C}^n(G)$ is connected unless G is a path (and $n > 1$) or G is a cycle (and $n > 2$). In all cases, however, the inequality $|G| \geq n$ ensures each connected component of $\mathcal{C}^n(G)$ contains a 0-cell of $\mathcal{D}^n(G)$. Thus, the inclusion $i : \mathcal{D}^n(G) \to \mathcal{C}^n(G)$ induces a surjection on π_0.

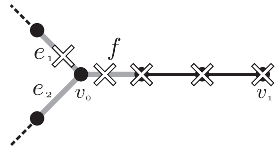

Figure 2: In the figure, each X represents a token. In $\mathcal{C}^n(G)$, the two tokens on edges e_1 and f can exchange places in a "dance" at the vertex v_0, using e_1, e_2, and f, while tokens occupy the remaining vertices on the essential path γ with endpoints v_0, v_1. However, no such dance is possible in $\mathcal{D}^n(G)$.

Next, Lemma 2.1 implies that i induces an injection on π_0 and, taking as α any loop based at a point of $\mathcal{D}^n(G)$, a surjection on π_1. Lemma 2.2 implies that i induces an injection on π_1. Since the connected components of $\mathcal{C}^n(G)$ are aspherical [9], as are those of $\mathcal{D}^n(G)$ [1], the Whitehead theorem establishes that i is a homotopy equivalence. Lastly, the pair $(\mathcal{C}^n(G), \mathcal{D}^n(G))$ satisfies the homotopy extension property, so a homotopy inverse to i extends to a deformation retraction $\mathcal{C}^n(G) \to \mathcal{D}^n(G)$.

$(\Rightarrow)$

The existence of a deformation retraction $\mathcal{C}^n(G) \to \mathcal{D}^n(G)$ implies the inclusion-induced homomorphism $i_* : \pi_1(\mathcal{D}^n(G)) \to \pi_1(\mathcal{C}^n(G))$ is an isomorphism. We show that

(i) if (A) is violated, then i_* either fails to be surjective or fails to be injective, and
(ii) if (B) is violated, then i_* fails to be surjective.

Abrams gives the following succinct argument for (ii). Suppose G is a graph for which (B) does not hold. Let λ be a cycle in G, which is of length at most n. Then the homotopy class of the loop in $\mathcal{C}^n(G)$ which simultaneously moves all n tokens around λ is a nontrivial element of $\pi_1(\mathcal{C}^n(G))$ which has no preimage under i_*, since at least $n+1$ edges are needed just to move a single token along the cycle for configurations in the discretized space.

For (i), we consider two separate cases. An *essential path* in a graph is a connected subgraph with two distinct essential vertices as endpoints, and no essential vertices in its interior. Let γ be an essential path in G with endpoints v_0, v_1, and assume γ is of length $m < n-1$, so that (A) does not hold. Let $h_1, h_2, \ldots, h_m$ be the edges of γ, labeled from v_0 to v_1.

<u>Case 1:</u> γ has an endpoint v_1 of degree 1.

If (v_0) is also of degree 1, then $|G| = |\gamma| < n$ (using the assumption G is connected), which implies $\mathcal{D}^n(G) = \emptyset$. So in this case, the other essential endpoint v_0 of γ is necessarily of degree ≥ 3. Let e_1, e_2 be edges of G incident with v_0, but not contained in γ, and let f be the edge of γ incident with v_0 (see Figure 2). Consider the loop $\alpha \subset \mathcal{C}^n(G)$ which fixes m tokens on the vertices of $\gamma \setminus \{v_0\}$, while two tokens trade places, then trade places again between v_0 and the other endpoint of e_1 – by a "dance" at v_0 using e_1, e_2,

and f – and the remaining $n-m-2$ tokens are fixed at any other vertices of G. (We mean a nontrivial pure braid which is similar to the square of the braid illustrated in Figure 1.) There is no homotopy class of $\pi_1(\mathcal{D}^n(G))$ which maps to the class $[\alpha]$ under i_*, since there are too few vertices in γ to execute this particular dance at v_0 in $\mathcal{D}^n(G)$. Thus, i_* is not surjective, so we conclude there is no deformation retraction $\mathcal{C}^n(G) \to \mathcal{D}^n(G)$.

Case 2: both endpoints of γ have degree ≥ 3.

In this case, we show the existence of a nontrivial loop $\alpha \subset \mathcal{D}^n(G)$ whose image under i is a contractible loop in $\mathcal{C}^n(G)$. This implies i_* has a nontrivial kernel, and therefore is not injective.

Let G' be the graph obtained from G by deleting $\mathrm{int}(\gamma)$. Let e_1, e_2 be edges of G', such that $e_1 \cap e_2 = v_0$; let w_1, w_2, respectively, be the other endpoints of e_1 and e_2. We may assume $w_1 \neq w_2$, since, if $w_1 = w_2$, then (B) is violated. Let f_1, f_2 be edges of G', both incident with v_1, whose other endpoints we label z_1 and z_2, respectively; as before, $z_1 \neq z_2$. We may further assume $w_i \neq z_j$ for $i, j \in \{1, 2\}$, since identifying any of these pairs of points would give a cycle $\gamma \cup e_i \cup f_j$ of length $m+2 \leq n$, violating (B). Let H be the subgraph $\gamma \cup \{e_1, e_2, f_1, f_2\}$. See Figure 3.

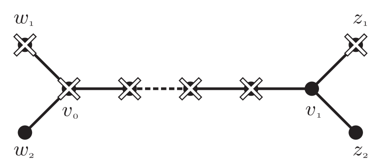

Figure 3: The subgraph $H \subset G$ in Case 2, where a nontrivial loop in $\mathcal{D}^n(G)$ becomes trivial under inclusion into $\mathcal{C}^n(G)$. The configuration illustrated is Step 1 of the loop described in Case 2.

The construction below assumes the existence of configurations in which $n-m-2$ tokens lie on vertices of $G \setminus H$, the remaining $m+2$ tokens lie on vertices of H, and three vertices of H are unoccupied. The existence of these configurations is not guaranteed by the assertion, proved earlier, that G contains at least n vertices. However, we may assume $n < |G| - 2$, as the following observations will show. Suppose to the contrary that $n \geq |G| - 2$. If G contains a cycle λ, then at least two vertices of H are not contained in λ, so condition (B) clearly does not hold. On the other hand, if G is a tree, then it has a vertex of degree 1, which is an endpoint of an essential path σ not containing at least four vertices of H. An application of Case 1 to σ then shows i_* is not surjective.

Let α be a non-contractible loop in $\mathcal{D}^n(G)$ which proceeds through the following numbered steps. The tokens $T_{m+3}, \ldots, T_n$ are fixed at vertices of $G \setminus H$ throughout. Here we use our assumption that $n < |G| - 2$, leaving three vertices of H unoccupied in Step 1 (the vertices v_1, w_2, and z_2, as illustrated in Figure 3) and at the end of each movement of a token along an edge in Steps 2–9 below.

1) Tokens T_1 and T_2 are at w_1 and z_1, respectively; tokens $T_3, \dots, T_{m+2}$ are on the vertices of $\gamma \setminus \{v_1\}$.
2) $T_{m+2}, T_{m+1} \dots, T_3$ each shifts one edge in the direction of v_1, after which these m tokens occupy the vertices of $\gamma \setminus \{v_0\}$.
3) T_1 moves from w_1 to v_0 to w_2.
4) $T_3, \dots, T_{m+2}$ shift back to $\gamma \setminus \{v_1\}$.
5) T_2 moves from z_1 to v_1 to z_2.
6) $T_{m+2}, T_{m+1} \dots, T_3$ shift back to $\gamma \setminus \{v_0\}$.
7) T_1 moves from w_2 to v_0 to w_1.
8) $T_3, \dots, T_{m+2}$ shift back to $\gamma \setminus \{v_1\}$.
9) T_2 moves from z_2 to v_1 to z_1.

In $\mathcal{C}^n(G)$, $i(\alpha)$ is homotopic to a loop α' which fixes tokens $T_3, \dots, T_{m+2}$ on the midpoints of $h_1, h_2, \dots, h_m$, while T_1 and T_2 make the moves described in Steps 3, 5, 7, and 9. But α' is homotopic to a constant loop at the configuration in which T_1 and T_2 are fixed at w_1 and z_1, respectively, and, as in α', token T_j is fixed at the midpoint of h_{j-2} for $3 \leq j \leq m+2$. Therefore $i(\alpha)$ is contractible and $\ker i_* \neq 0$. □

3. Examples of optimal subdivision

The subdivision in a graph satisfying the conditions in Definition 1.1 is *sufficient* in the following sense. For any connected graph G, adding a finite number of non-essential vertices to G (that is, subdividing an edge e of G by removing the interior of e, and adding a vertex of degree 2 connected by an edge to each endpoint of e) gives a new graph G' which is homeomorphic to G. Thus, for fixed n, $\mathcal{C}^n(G)$ is homeomorphic to $\mathcal{C}^n(G')$. The deformation retraction whose existence is proved in Theorem 1 projects to the quotient $U\mathcal{C}^n(G) \to U\mathcal{D}^n(G)$, so the theorem holds for unordered spaces also. Although $U\mathcal{D}^n(G)$ and $U\mathcal{D}^n(G')$ may be very different spaces in a combinatorial sense, the theorem guarantees that if $\pi_1(U\mathcal{D}^n(G))$ is isomorphic to $B_n(G)$, then $\pi_1(U\mathcal{D}^n(G'))$ is isomorphic to $B_n(G)$, too.

In the following examples, we compute 3-strand braid groups of two graphs, each optimally subdivided for 3 strands, using unlabeled discretized configuration spaces.

Example 3.1
Consider the graph G, at left in Figure 4, homeomorphic to the letter P, where the cycle contains four edges and the essential path contains two edges. It is evident from these specifications that G is optimally subdivided for $n = 3$ strands. That is, the fundamental group of $U\mathcal{D}^3(G)$ is isomorphic to $B_3(G)$, as is $\pi_1(U\mathcal{D}^3(G'))$, if G' is any graph obtained from G by subdividing edges of G; but the discretization fails with any fewer edges in the cycle or fewer edges in the essential path.

We illustrate $U\mathcal{D}^3(G)$ at right in Figure 4. For a cubical complex X of dimension r, let the *maximal-cell vector* of X be the vector $c = (c_1, \dots, c_r)$, where the positive integer c_k is the number of top-dimensional k-cells in X. As the figure shows, the maximal-cell vector of $U\mathcal{D}^3(G)$ is $(4, 4, 2)$. Each 3-cell is the product of three disjoint edges; they are $\alpha \times \epsilon \times \gamma$ and $\alpha \times \delta \times \phi$. Each top-dimensional 2-cell is the product of two edges and one vertex; they

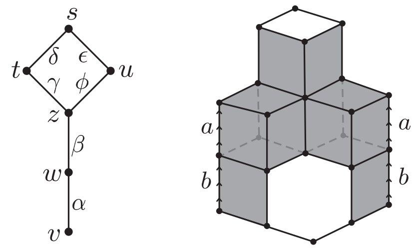

Figure 4: At left is the graph G, optimally subdivided for 3 strands. At right is the discretized configuration space $U\mathcal{D}^3(G)$.

are $\beta \times \delta \times u$, $\beta \times \epsilon \times t$, $\beta \times \delta \times v$, and $\beta \times \epsilon \times v$. Each top-dimensional 1-cell is the product of two vertices and a single edge. They are $s \times u \times \gamma$, $s \times t \times \phi$, $v \times w \times \gamma$, and $v \times w \times \phi$.

Although this space can be embedded in $\mathbb{R}^3$ with some distortion, the realization illustrated here has the advantage that every k-cell appears as a k-cube with uniform side length. The ambient space can be thought of as $\mathbb{R}^2 \times S^1$, obtained by an orientation-preserving quotient of $\mathbb{R}^3$ under which points of two planes tangent to the configuration space along $\{a \cup b\}$ are equivalent. Hence, the 1-cells labeled a (resp. b) are identified, with the orientation shown in the figure. The configuration space has an automorphism of order 2, not readily apparent in the illustration, which, for example, switches the 0-cells at the top and bottom of the picture. These 0-cells are the configurations $\{v, w, z\}$ and $\{s, t, u\}$. If we take the 0-cell at the bottom to be $\{v, w, z\}$, then the 1-cell labeled a in the figure is $\alpha \times s \times z$, and b is $\beta \times s \times v$.

Note the 4-cycle at the top of the figure is non-trivial, as is the 4-cycle consisting of the four 1-cells at the bottom. The space deformation retracts to the wedge sum of three circles; hence, $B_3(G)$ is isomorphic to F_3, the free group of rank 3.

Example 3.2
Now consider the graph Γ, at left in Figure 5. It is homeomorphic to the letter Y, with one vertex of degree 2 on each of the arms. It, too, is optimally subdivided for 3 strands. Its configuration space $U\mathcal{D}^3(\Gamma)$, at right in the figure, has a nice property one could call "unit lattice embedding" in $\mathbb{R}^3$. That is, it can be embedded in $\mathbb{R}^3$ such that every k-cell is a k-cube in $\mathbb{R}^3$ with side length 1, whose vertices lie in $\mathbb{Z}^3$. The maximal-cell vector of $U\mathcal{D}^3(\Gamma)$ is (6,6,4). The 3-cell which shares a 2-cell face with each of the other three 3-cells is the product of edges $\alpha \times \delta \times \epsilon$. We leave the product descriptions of the other top-dimensional cells as an exercise for the interested reader, here only remarking that $\mathrm{Aut}(\Gamma)$ is nontrivial, and consequently, various labelings are possible. As in the previous example, the space deformation retracts to the wedge sum of three circles, so $B_3(\Gamma)$ too is isomorphic to F_3.

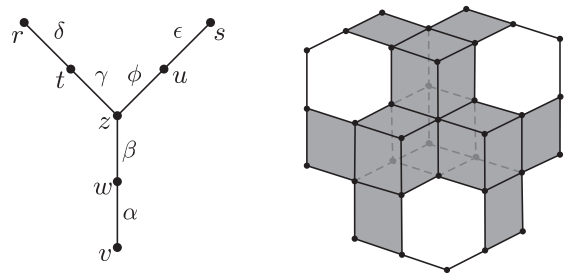

Figure 5: At left is the graph Γ, optimally subdivided for 3 strands. At right is the discretized configuration space $U\mathcal{D}^3(\Gamma)$.

4. Paths and homotopy in $\mathcal{D}^n(G)$

The statements of Lemmas 2.1 and 2.2 below are identical to those of Lemmas 2.2 and 2.3 of [1]. We need to slightly modify the proof of each lemma, given the relaxed hypothesis (A) in Theorem 1. The substance of the modifications is in much more careful tracking of the correspondence between 0-cells of $\mathcal{D}^n(G)$ and certain configurations in $\mathcal{C}^n(G) \setminus \mathcal{D}^n(G)$, and the standardized moves in $\mathcal{D}^n(G)$ between these 0-cells. In the statement of Lemma 2.2, we write D^2 to denote a 2-dimensional disk.

Lemma 2.1 (Path Homotopy Lemma). Let $\alpha : (I, \partial I) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$. If conditions (A) and (B) of Theorem 1 hold then α can be homotoped rel ∂I into $\mathcal{D}^n(G)$.

Lemma 2.2 (Contractible Loop Lemma). Let $f : (D^2, \partial D^2) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$. If conditions (A) and (B) of Theorem 1 hold then there exists $g : D^2 \to \mathcal{D}^n(G)$ such that $g|_{\partial D^2} = f|_{\partial D^2}$.

Let d be the standard metric on G induced by identifying each edge homeomorphically with the unit interval. Let $0 < \delta < 1/2$, and let $S_\delta = \{y \in G \mid d(y, v) = \delta \text{ for some essential vertex } v\}$. Let the connected components of $G \setminus S_\delta$ be labeled $X_1, X_2, \dots, X_r$. The connected components containing essential vertices of G are *vertex components*, while the rest are termed *arc components*. We give an arbitrary orientation on each arc component, and say the orientation is directed away from the *tail* of the arc, and toward its *head*.

Next, suppose $x \in \mathcal{C}^n(G)$ is a configuration with at most one token in each vertex component, and no tokens on points of B_δ. We define the *combinatorics of* x to be the collection $K(x) = (\kappa_1(x), \kappa_2(x), \dots, \kappa_r(x))$, where $\kappa_j(x)$ is an ordered list of tokens of x lying on X_j, and we say x is a *configuration with combinatorics*. For an arc component X_j, the ordering in the list κ_j comes from the orientation on X_j. By our choice of x, no ordering is needed for the vertex components, since each contains at most one token.

The partition of G into $S_\delta, X_1, \dots, X_r$ extends naturally to a partition of $\mathcal{C}^n(G)$, as follows. For each $1 \le j \le n$, let $p_j : \mathcal{C}^n(G) \to G$ be the projection

onto the j-th coordinate. Let

$$S = \bigcup_j p_j^{-1}(S_\delta).$$

Then $\mathcal{C}^n(G)$ is the disjoint union of S and its complement S^c, where, if x is a configuration with combinatorics in some connected component $\mathcal{L}_x$ of S^c, then $K(y) = K(x)$ for every configuration $y \in \mathcal{L}_x$. We note that elements of S do not have combinatorics, nor does any configuration $x \in \mathcal{C}^n(G)$ where two or more points of x are in the δ-neighborhood of the same essential vertex of G.

Let us prove two assertions about this partition of $\mathcal{C}^n(G)$.

(i) For each x with combinatorics, the set

$$\mathcal{L}_x = \{y \in \mathcal{C}^n(G) \mid x, y \text{ have the same combinatorics}\}$$

is contractible.

Proof: $\mathcal{L}_x$ is a product of possibly some 1-strand configuration spaces of vertex components, and possibly some multiple-strand configuration spaces of arc components. Both types of factors are contractible; hence, so is the product $\mathcal{L}_x$.

(ii) If x and y have the same combinatorics, then y is in the same connected component of S^c as x.

This follows from (i).

Assuming conditions (A) and (B) of Theorem 1 hold, to each configuration in $\mathcal{C}^n(G)$ with combinatorics, we associate exactly one 0-cell of $\mathcal{D}^n(G)$ called a *standard vertex*. (The definition is below, numbered 4.1.) Tokens of a configuration x with combinatorics are in *standard position* if, for each arc component X_j containing m tokens, the tokens lie on the first m vertices of X_j, according to the orientation on the component. We express the notion of standard vertex in terms of a function – one which is not necessarily injective. Indeed, we cannot require standard vertices to have the same combinatorics as the configurations to which they are assigned, since under condition (A), arc components may contain $n-2$ or $n-1$ vertices. Instead, we define the standard vertex as follows, treating the cases $n > 2$ and $n = 2$ separately.

Definition 4.1 (Standard vertex)**.** First, suppose $n > 2$. (We consider $n = 2$ below.) Let $x \in \mathcal{C}^n(G)$ be a configuration with combinatorics. For an arc component X_i, let q_i denote the number of vertices of G lying in X_i. Recall that $|\kappa_i(x)|$ is the number of tokens of x on the component X_i. When G is sufficiently subdivided, note that $q_i \geq n-2$. Then the quantity $\theta_i(x) = |\kappa_i(x)| - q_i$ measures the "overcrowding" of tokens on X_i. If $\theta_i(x) \leq 0$ for all i, the *standard vertex associated with* x, denoted $\Phi(x)$, is simply the vertex of $\mathcal{D}^n(G)$ with the same combinatorics as x, such that for each arc component X_j, the tokens in X_j (if any) are in standard position. Sufficient subdivision ensures that $\theta_i < 3$ for all i. The cases $\theta_i = 1, 2$ may occur on essential paths of length less than $n+1$; these situations require further elaboration.

Let γ be the essential path in G containing the arc component X_i, and give the labels w and z to the endpoints of γ, so that X_i is oriented toward z. Denote the vertex components containing w and z by X_w and X_z, respectively.

Case 1: $\theta_i(x) = 1$.

We first note that at most one token lies in some component different from X_i. Since $\theta_i(x) = 1$, either $q_i = n-2$ and $|\kappa_i(x)| = n-1$, or $q_i = n-1$ and $|\kappa_i(x)| = n$. Using the orientation on X_i, let us label the first $n-1$ tokens lying in X_i by $T_1, T_2, \dots, T_{n-1}$, and give the label T_n to the remaining token. The standard vertex $\Phi(x)$ depends on the location of the component X_j containing T_n.

(a) If $X_j = X_w$, then $\Phi(x)$ is the standard vertex whose combinatorics for X_w, X_i, X_z are, respectively, $(T_n), (T_1, \dots, T_{n-2}), (T_{n-1})$.
(b) If $X_j = X_i$, then $q_i = n-1$ and $\Phi(x)$ is the standard vertex whose combinatorics for X_w, X_i are $(T_1), (T_2, \dots, T_n)$.
(c) If $X_j \notin \{X_w, X_i\}$, then $\Phi(x)$ is the standard vertex whose combinatorics for X_w, X_i, X_j are $(T_1), (T_2, \dots, T_{n-1}), (T_n)$.

Case 2: $\theta_i(x) = 2$.

For sufficiently subdivided graphs, $\theta_i(x) = 2$ implies $q_i = n-2$ and $|\kappa_i(x)| = n$. We label the tokens $T_1, T_2, \dots, T_n$ according to the orientation on X_i. The standard vertex $\Phi(x)$ is the standard vertex whose combinatorics for X_w, X_i, X_z are $(T_1), (T_2, \dots, T_{n-1}), (T_n)$.

One small modification is needed when $n = 2$, since essential paths of G need only contain one edge. Suppose two essential paths, each of length 1, share a common endpoint w. If both arc components are oriented so that w is at the tail, and the two tokens of a configuration x occupy both arc components, there is a problem: both tokens of $\Phi(x)$ cannot simultaneously lie on w. For simplicity, we make the following definition, regardless of the orientations on the arc components. Let e be an edge of G with endpoints the essential vertices w and u; let f be an edge whose endpoints are w and another essential vertex z. We observe that, by condition (B) of the theorem, $u \neq z$. Denote the corresponding arc components of $G \setminus S_\delta$ by X_e and X_f, and the vertex components by X_u, X_w, X_z. If x has combinatorics $(T_1), (T_2)$ on the arc components X_e, X_f, then $\Phi(x)$ is the vertex of $\mathcal{D}^n(G)$ with combinatorics $(T_1), (T_2)$ on the vertex components X_u, X_z.

The last preliminary step before we give the proofs of Lemmas 1 and 2 is to consider the moves between standard vertices – paths in the discretized configuration space corresponding to the movement of a token from a vertex component to an arc component, or vice-versa. We call these *standard moves.* We begin by restating the standard moves between configurations with non-positive values of θ_i, as in [1]. As above, let γ be an essential path with endpoints w and z, containing an arc component X_i which is oriented toward z. If x is a configuration with combinatorics

- $(T_1), (T_2, \dots, T_k)$ on the components X_w, X_i, or

- $(T_1, \ldots, T_{k-1}), (T_k)$ on the components X_i, X_z,

when a token enters X_i from the vertex component X_w resp. X_z, the combinatorics of the standard vertex associated with the new configuration x' are $\kappa_i(\Phi(x')) = (T_1, T_2, \ldots, T_{k-1}, T_k)$ and $|\kappa_w(x')| = 0$ resp. $|\kappa_z(x')| = 0$.

Now we consider the remaining moves, which begin and/or terminate at a standard vertex $\Phi(x)$ with $\theta_i(x) > 0$ for some i. Let x be a configuration such that $\theta_i(x) > 0$ for some arc component X_i. Let γ be the essential path that contains X_i. Let w and z be the endpoints of γ, so that X_i is oriented from w toward z. Again, let X_w and X_z denote the vertex components containing w and z, respectively. The standard moves below are defined for general values of n. In the case $n = 2$, they make sense, except that no tokens of the configuration represented by a standard vertex lie in X_i if X_i contains no vertices of G, and the standard move numbered 4 does not occur.

Each case is illustrated with a figure. Each figure shows an essential path, its endpoints circled to suggest the points of S_δ. Each X represents a token in the relevant configurations in $\mathcal{C}^n(G)$; the bolded X is the token moving between an arc component and a vertex component. Orientation on the arc components is illustrated by the symbols $\rangle$ and $\langle$. The upper portion of each figure shows the corresponding standard configuration in $\mathcal{D}^n(G)$, with tokens placed at the vertices represented by filled-in circles.

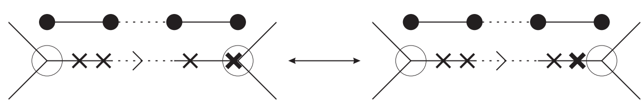

Figure 6: The two configurations depicted in this figure illustrate movement of a token between components of $G \setminus S_\delta$. Above each configuration, the associated standard vertex is depicted. Here, token T_n, illustrated with the bold X, enters (or leaves) the head of a minimally subdivided arc component containing the other $n-1$ tokens. The corresponding (trivial) move in $\mathcal{D}^n(G)$ is the constant path at the single standard vertex associated with both configurations.

(1) Suppose $\theta_i(x) = 1$, and $\kappa_i(x) = (T_1, \ldots, T_{n-1})$, while $\kappa_z(x) = (T_n)$ and $|\kappa_w(x)| = 0$. Let token T_n enter X_i from the head. The corresponding standard move is trivial, since the "before" and "after" configurations have the same combinatorics on X_w, X_i, and X_z. See Figure 6.
(2) Suppose $\theta_i(x) \in \{0, 1\}$, and $\kappa_i(x) = (T_1, \ldots, T_{n-2})$ (and possibly T_{n-1}), while $\kappa_w(x) = (T_n)$ and $|\kappa_z(x)| = 0$. Let token T_n enter X_i from the tail. Again the corresponding standard move is trivial, with fixed combinatorics on X_0, X_i, and X_1. See Figure 7.
(3) Suppose $\theta_i(x) = 1$, and $\kappa_i(x) = (T_1, \ldots, T_{n-1})$, while $\kappa_w(x) = (T_n)$ and $|\kappa_z(x)| = 0$. Let token T_n move from X_w to an arc component $X_j, j \neq i$. Then the standard move takes T_n to standard position on X_j, while moving T_1 to the vertex component X_w, and placing the remaining tokens in standard position on X_i. See Figure 8.

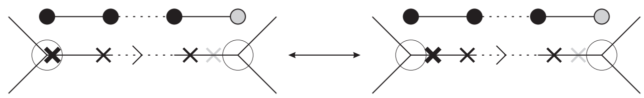

Figure 7: Token T_n enters (or leaves) the tail of an arc component that is already "full," i.e., with $\theta_i \geq 0$. Token T_{n-1} is shaded in grey, to indicate that the associated standard vertex does not depend on the presence of T_{n-1}.

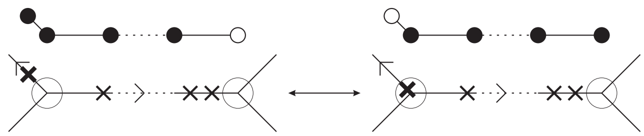

Figure 8: Token T_n enters (or leaves) X_0, with X_i containing the other $n-1$ tokens.

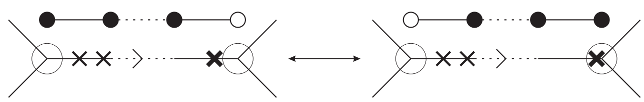

Figure 9: Token T_{n-1} leaves (resp. enters) X_i, after which $\theta_i(x) = 0$ (resp. $\theta_i(x) = 1$).

(4) Suppose $\theta_i(x) = 1$, and $\kappa_i(x) = (T_1, \ldots, T_{n-1})$, while $|\kappa_w(x)| = |\kappa_z(x)| = 0$. Let T_{n-1} leave X_i from the head, entering X_z. Then the standard move is the path in $\mathcal{D}^n(G)$ that takes T_1 from the first vertex of X_i to z, shifts each T_{p+1} to the p-th vertex of X_i for $p = 1, 2, \ldots, n-3$, and moves T_{n-1} from w to vertex $n-2$ of X_i. See Figure 9.

Now that we have completed all of the preliminary steps, we are now ready to proceed with our proofs of Lemma 2.1 and Lemma 2.2.

Proof of Lemma 2.1. We construct a path $\beta : I \to \mathcal{D}^n(G)$ which is homotopic rel boundary to a given path $\alpha : (I, \partial I) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$.

Pick $\delta < 1/4$ small enough that for every configuration $x \in \alpha(I)$, the distance between any pair of tokens is at least 4δ. Since each point of B_δ has a manifold neighborhood in G, we can perturb α so that it is transverse to the set S, with the pairwise distance between tokens bounded below by 3δ after the perturbation. We call the path thus obtained α as well. Observe that each $x \in \alpha(I) \backslash S$ has combinatorics, since two tokens in the same vertex component are strictly less than 2δ apart; our choice of δ precludes such a configuration in $\alpha(I)$.

Since α is transverse to S, the set $\alpha^{-1}(S)$ is a finite subset of I. Let us label its points $t_1, t_2, \ldots, t_m$. Let $t_0 = 0$ and let $t_{m+1} = 1$. For each index $1 \leq i \leq m+1$, arbitrarily choose $t_i^* \in (t_{i-1}, t_i) \subseteq I$. The points

$\alpha(t_i^*)$ have combinatorics, by the observation in the preceding paragraph. For each t_i^*, let v_i^* be the standard vertex in $\mathcal{D}^n(G)$ associated with $\alpha(t_i^*)$. With the standard moves defined above, we have shown that consecutive configurations v_i^* can be joined by edge paths in $\mathcal{D}^n(G)$.

We construct β as follows. Let β_0 be a path from $\alpha(0)$ to v_1^* which stays in $\mathcal{D}^n(G)$. (For example, see the construction in [1].) For $1 \leq i \leq m$, let β_i be the path in $\mathcal{D}^n(G)$ from v_i^* to v_{i+1}^* given by the appropriate standard move. Let β_{m+1} be a path from v_{m+1}^* to $\alpha(1)$ which stays in $\mathcal{D}^n(G)$. Let β be the concatenation of the paths $\beta_0 \cdot \beta_1 \cdots \beta_{m+1}$.

Now we show $\beta \simeq \alpha$. Reparameterize β so that $\beta^{-1}(S) = \alpha^{-1}(S)$. For $t \in I$, if $x = \alpha(t)$ has combinatorics, then $\beta(t)$ has the same combinatorics. Since $\mathcal{L}_x$ is contractible by fact (i) above, the arc $\alpha((t_{j-1}, t_j))$ is homotopic to $\beta((t_{j-1}, t_j))$ for each $j = 1, 2, \ldots, m+1$, by a linear homotopy. On the other hand, for $t \in \alpha^{-1}(S)$, we homotope between $\alpha(t)$ and $\beta(t)$ by a linear homotopy restricted to the tokens not lying in B_δ. □

Proof of Lemma 2.2. Given a map $f : (D^2, \partial D^2) \to (\mathcal{C}^n(G), \mathcal{D}^n(G))$, we construct another map $g : D^2 \to \mathcal{D}^n(G)$ which agrees with f on the boundary of the disk.

As in the proof of Lemma 2.1, we choose δ small enough that in every configuration $x \in f(D^2)$, the n tokens lie on points of G whose pairwise distance is at least 4δ. We preserve a lower bound of 3δ after perturbing f so that its image is transverse to S. This ensures that every point in $f(D^2) \setminus S$ is a configuration with at most one token in each vertex component, and therefore has combinatorics.

We consider the set $Z := f^{-1}(S) \subset D^2$. By transversality, it is a union of circles and proper arcs in D^2. Each circle and each proper arc is the inverse image of an intersection of $f(D^2)$ with an $(n-1)$-flat in $\mathcal{C}^n(G)$; each can be labeled by a single point of B_δ and a single token. Observe that f maps points in the same connected component of $D^2 \setminus Z$ to configurations that have the same combinatorics, since a transverse intersection of $f(D^2)$ with S corresponds to a token moving between an arc component and a vertex component of G.

To simplify the process of describing the map g based on Z, we remove connected components of Z that do not intersect ∂D^2. Suppose Y is such a component. Let $\Delta \subset D^2$ be a closed (topological) disk with boundary circle C, such that $Y \subset \Delta$ and $C \cap Z = \emptyset$. Let $c \in C$, and observe that all points of $f(C)$ have the same combinatorics as $f(c)$. We define a new map $\hat{f}$ to agree with f on $D^2 \setminus \Delta$, and to map Δ to the cone from the set $f(C)$ to the standard vertex $\Phi(c)$.

After all components of Z have been removed except those that touch ∂D^2, we call the resulting map f also. Let $H \subset D^2$ be a "dual graph" of Z, constructed as follows. Place a vertex of H in each component of $\mathrm{int}(D^2) \setminus Z$. Also place a vertex of H in each arc of $\partial D^2 \setminus Z$. Connect two vertices in ∂D^2 by an edge that is a subarc of ∂D^2 if that edge can be made to intersect Z in only one point. Connect a vertex $w_1 \in \partial D^2$ to a vertex $w_2 \in \mathrm{int}\,(D^2)$ if the two points lie in the same connected component of $D^2 \setminus Z$. Lastly, if

two vertices $w_1, w_2 \in \operatorname{int}(D^2)$ can be connected by an edge intersecting Z in exactly one point, we do so (see Figure 10).

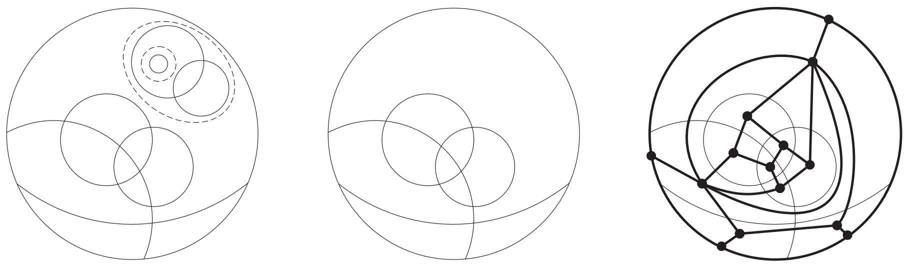

Figure 10: We first remove non-essential components of Z, then form the dual graph H, thus simplifying the description of the map g.

Let us analyze the structure of H. There are two types of edges in H: those that have a point of intersection with Z (called *Z-intersecting* edges), and those that do not (called *Z-disjoint* edges). The former are edges between two vertices, both in ∂D^2 or both in $\operatorname{int}(D^2)$; the latter are edges between one vertex in the boundary and one vertex in the interior of D^2. A region in $D^2 \setminus H$ is a quadrilateral, and may be classified as follows:

- a *boundary region* is bounded by two Z-disjoint edges intersecting the same arc of Z, and two opposite edges;
- an *interior region* is bounded by two pairs of opposite edges, each pair intersecting a distinct arc of Z.

To see that the regions are quadrilaterals, observe that the closure of each region either contains a point of intersection of two arcs of Z or contains a point of intersection of Z with ∂D^2. In the former case, opposite edges of the boundary of the region are Z-intersecting edges, each intersecting the same arc of Z; in the latter case, the boundary of the region contains one such pair (with one edge lying in ∂D^2), while each of the other two is a Z-disjoint edge.

Now we can describe the map $g : D^2 \to \mathcal{D}^n(G)$. First, we map the edges of $H \subset D^2$, according to type. If e is a Z-disjoint edge, then it is mapped to the single standard vertex of $\mathcal{D}^n(G)$ associated with $f(x)$ for every $x \in e$. If e is Z-intersecting, with endpoints x_0 and x_1, then we map e using the standard move that connects the standard vertices associated with $f(x_0)$ and $f(x_1)$.

Having mapped edges of H, we extend the map over the interior of each quadrilateral region. For a boundary region R, g factors through a fibration with total space R and base space the Z-intersecting edge. For an interior region Q, we first observe that Q contains two subarcs γ_1 and γ_2 of Z, which intersect transversely in Q. Thus, γ_1 and γ_2 are labeled by different tokens. To see that the points of B_δ labeling γ_1 and γ_2 are distinct, observe that a traversal of $f(\gamma_1)$ corresponds to keeping some token T_1 fixed at a point $b_1 \in B_\delta$ while (as the path crosses $f(\gamma_2)$) some token $T_2, k_1 \neq k_2$, crosses a point $b_2 \in B_\delta$. Since $f(\gamma_1 \cap \gamma_2)$ is a point in $\mathcal{C}^n(G)$ where no two tokens are within 3δ of each other, the points b_1 and b_2 cannot be within δ of the

same essential vertex of G. Thus, the two points are distinct, but we obtain an even stronger condition: they do not lie in the closure of a single vertex component. In the typical case, this gives the following property:

(*) *the edges used in the standard move to cross γ_1 are disjoint from the edges used in the standard move to cross γ_2.*

Thus, g extends over Q as a map from D^2 to $\mathcal{D}^n(G)$, simultaneously homotoping between the standard vertices on opposite sides of γ_1 while homotoping between the standard vertices on opposite sides of γ_2.

However, the property labeled (*) above does not hold in certain cases. Let U be the set of four vertices of H contained in the boundary of Q. Let μ be an essential path in G with endpoints w and z. We write X_m to denote the arc component contained in μ, and X_w (resp. X_z) to denote the vertex component containing w (resp. z). Let b_1 and b_2 be points of B_δ located at distance δ from w and z, respectively, and labeling the arcs γ_1 and γ_2, respectively. Property (*) holds in all except the following cases.

(a) The points b_1 and b_2 are the endpoints of the closure of X_m, and, for each $u \in U$, at most $q_m + 1$ tokens of $f(u)$ lie on X_m.
(b) Exactly one of the two points of B_δ lies in μ, and $q_m + 2$ tokens of $f(u)$ lie in $X_m \cup X_w \cup X_z$ for some $u \in U$.

Case (b) does not arise in graphs with essential paths subdivided to at least $n+1$ edges, as in the hypotheses of Theorem A. However, in the proof of Lemma 2.3 of [1], case (a) is not considered, which is a slight omission we remedy below. In both cases, we must explicitly define the extension of g over the interior of Q.

Figures 11, 12, and 13 each illustrates such an extension. The interior region $Q \subset D^2$ is represented by the square in the center of each figure. This square is decorated in each figure, in a way that will be explained below. The two graph diagrams at each corner of the square represent a configuration in $\mathcal{C}^n(G)$ (the upper graph, with tokens denoted X, and B_δ suggested by the dashed circles around essential vertices) and the associated standard vertex in $\mathcal{D}^n(G)$ (lower graph, with filled circles denoting occupied vertices). Each edge of Q is an edge of the dual graph H, and thus, is mapped to the standard move between the images of its own endpoints, as observed previously.

We first suppose $q_m > 0$. (When $n = 2$, some arc components of $G \setminus B_\delta$ may contain no vertices of G; we address this situation below.) Figure 11 represents an extension of g over Q when none of the standard moves in $g(\partial Q)$ is trivial. This occurs in case (a), provided $\theta_m(x) \leq 0$ for all $x \in g(\partial Q)$. Here, we define the extension such that $g|_Q$ is one factor in a composite homotopy rel endpoints between two paths in $\mathcal{D}^n(G)$. This assertion will be made more clear below. First, label the vertices of the graph H lying on ∂Q as follows. Let v_0 be the vertex such that tokens of the configuration $f(v_0)$ occupy both X_w and X_z; let v_1 be the opposite vertex, such that tokens of $f(v_1)$ occupy neither vertex component. Arbitrarily label the other two vertices u_0 and u_1. Let p_0 be a path starting at v_0 and traversing two edges of Q, going through u_0 to v_1. Let p_1 be a similar path across the other two edges of Q, starting at v_0, going through u_1, and ending at v_1. Then $g(p_0) \simeq g(p_1)$ rel $\{g(v_0), g(v_1)\}$, via a homotopy remaining in $\mathcal{D}^n(G)$.

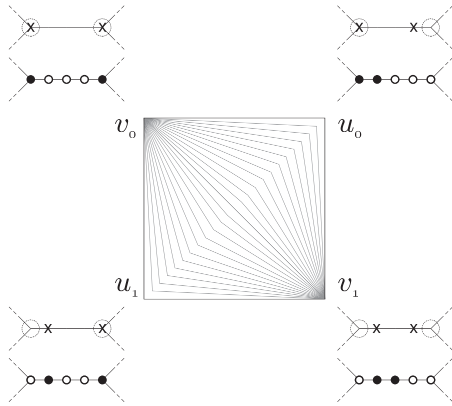

Figure 11: The extension of g over the interior region Q in case (a), involving a homotopy in $\mathcal{D}^n(G)$ between two paths using intersecting sets of edges of G.

To see how the image of this homotopy is also the image of Q under g, suppose $\mathcal{G} : I \times I \to \mathcal{D}^n(G)$ is a homotopy between the two paths. Further suppose $\mathcal{G} = \mathcal{F} \circ \mathcal{H}$, where $\mathcal{H} : I \times I \to Q$ maps segments $\{s\} \times I, s \in [0,1]$, to paths in Q from v_0 to v_1 such as those illustrated in Figure 11. Then $\mathcal{F} = g|_Q$.

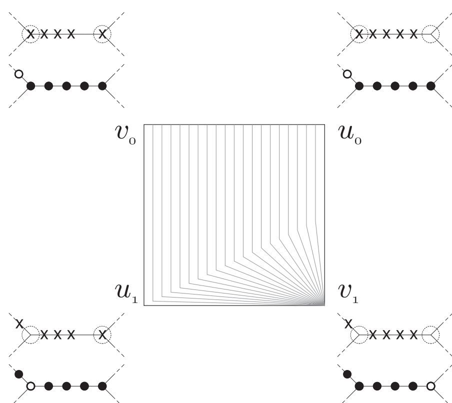

Figure 12: The extension of g over Q in case (b).

Figure 12 illustrates the extension of g over the interior of Q when exactly one of the standard moves $g(e)$, for e an edge of Q, is trivial. This type of extension is used both in case (a), when $-1 \leq \theta_i(x) \leq 1$ for all $x \in g(\partial Q)$, and in case (b), when the orientation on X_m is such that the standard move from $g(u_1)$ to $g(v_1)$, crossing b_2, is nontrivial. The latter situation is illustrated in Figure 12. Again, $g|_Q$ can be viewed as a composition factor of a homotopy. This time, however, the homotopy satisfies a stronger condition: it is the (continuous) reparameterization of a path. Again, let v_0 be the vertex of Q such that tokens of $f(v_0)$ occupy both vertex components

X_w and X_z, and again, let v_1 be the opposite vertex of Q, so that neither X_w nor X_z is occupied by tokens of $f(v_1)$. Label the other pair of opposite vertices u_0 and u_1, such that the standard vertices $\Phi(f(v_0)) =: g(v_0)$ and $\Phi(f(u_0)) =: g(u_0)$ have the same combinatorics. The edge with endpoints v_0 and u_0 (the top edge, in Figure 12) is then mapped by g to a trivial standard move. Let p_1 be a path in ∂Q from v_0 to u_1 to v_1, using the edges $\{v_0, u_1\}$ and $\{u_1, v_1\}$. Let p_0 be a path in ∂Q from u_0 to v_1 on the edge joining them (the right edge in Figure 12). Then $g(p_1) \simeq g(p_0)$, and again, the homotopy $I \times I \to \mathcal{D}^n(G)$ decomposes into two factors, one of which is $g|_Q$.

Finally, we address the cases that may occur when $q_m = 0$ (note this necessarily implies $n = 2$). Observe that this assumption precludes case (a): When $q_m = 0$, a configuration with both tokens on X_m exceeds the maximum $q_m + 1$ which defines case (a). (For any value of n, on an interior region Q where $\{b_1, b_2\} = \partial X_m$, but with $0 \leq \theta_m(u) \leq 2$ for each $u \in U$, each vertex of Q is mapped to the standard vertex $v \in \mathcal{D}^n(G)$ that has tokens occupying w and z. Each edge of Q is mapped by g to a trivial standard move. Thus, g is the constant map $g(x) = v$ for all x in the closure of Q.) It remains, then, to describe the extensions of g over Q in case (b). These map one or two edges of Q to a trivial standard move. Let λ be the essential path containing b_1. Let X_l be the arc component contained in λ. Let v_0 and v_1 be opposite vertices of Q, similar to the preceding descriptions and figures, such that $f(v_0)$ has combinatorics $(T_1), (T_2)$ on the vertex components X_w, X_z, and $f(v_1)$ has combinatorics $(T_1), (T_2)$ on the arc components X_l, X_m. Let u_1 be the vertex of Q such that $f(u_1)$ has combinatorics $(T_1), (T_2)$ on the components X_w and X_m. Since $q_m = 0$, this means the edge $\{v_0, u_1\}$ is always mapped by g to a trivial standard move in $\mathcal{D}^2(G)$.

Suppose first that $q_l > 0$. Then the edge $\{u_1, v_1\}$ is mapped to a nontrivial standard move precisely when the orientation on X_m is such that $\Phi(f(v_1))$ has combinatorics $(T_1), (T_2)$ on the components X_l, X_w. Since the edges $\{v_0, u_1\}$ and $\{u_0, v_1\}$ correspond to T_1 crossing b_1, and each is mapped to a nontrivial standard move no matter what orientation X_l is given, the orientation on X_m in the preceding sentence means there are no additional trivial standard moves in the image of ∂Q besides the image of $\{v_0, u_0\}$. In that case, the extension of g over Q is suggested by a diagram similar to that illustrated in Figure 12. In contrast, the opposite orientation on X_m makes the image of $\{u_1, v_1\}$ a second trivial standard move. The image of the opposite edge $\{v_0, u_0\}$ is also trivial, so g extends over the interior of Q with a product structure, similar to the mapping of a boundary region R.

For the remaining cases, with $q_l = 0$, the extension of g over Q depends on the orientation of X_l rather than that of X_m, since it is the orientation on X_l which determines the standard vertex $\Phi(f(u_1)) =: g(u_1)$ associated with configurations in which the tokens lie in X_z and X_l. Let y be the other endpoint of the single edge of the essential path λ. Recall from Definition 4.1 that, since $q_l = q_m = 0$, the standard vertex $\Phi(f(v_1)) =: g(v_1)$ does not depend on the orientations of X_l and X_m; it is the configuration (y, z), i.e., with token T_1 on the vertex y and token T_2 on the vertex z. If $g(u_1) = g(v_1)$, then g again maps the edge $\{u_1, v_1\}$ to a trivial standard move, meaning the

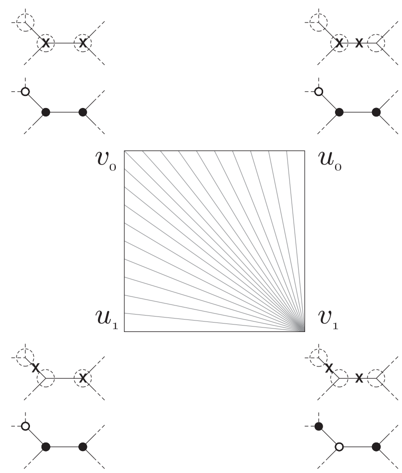

Figure 13: An extension of g over Q in case (b), when $n = 2$.

extension over Q is similar to the latter case in the preceding paragraph. The other possibility is that $g(u_1) = g(v_0)$, as illustrated in Figure 13. In this case, two adjacent edges of ∂Q (the top and left edges in the figure) are mapped by g to trivial standard moves. Then $g|_Q$ can again be viewed as a composition factor of a homotopy – this one fixing a parameterization of the path $g(\{u_0, v_1\}) = g(\{u_1, v_1\})$. □

References

[1] A. Abrams. *Configuration spaces and braid groups of graphs*. PhD thesis, UC Berkeley, 2000.

[2] Aaron Abrams and Robert Ghrist. Finding topology in a factory: configuration spaces. *Amer. Math. Monthly*, 109(2):140–150, 2002.

[3] Joan S. Birman and Tara E. Brendle. Braids: a survey. In *Handbook of knot theory*, pages 19–103. Elsevier B. V., Amsterdam, 2005.

[4] John Crisp and Bert Wiest. Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.*, 4:439–472, 2004.

[5] Daniel Farley. Homology of tree braid groups. In *Topological and asymptotic aspects of group theory*, volume 394 of *Contemp. Math.*, pages 101–112. Amer. Math. Soc., Providence, RI, 2006.

[6] Daniel Farley. Presentations for the cohomology rings of tree braid groups. In *Topology and robotics*, volume 438 of *Contemp. Math.*, pages 145–172. Amer. Math. Soc., Providence, RI, 2007.

[7] Daniel Farley and Lucas Sabalka. Discrete Morse theory and graph braid groups. *Algebr. Geom. Topol.*, 5:1075–1109 (electronic), 2005.

[8] Daniel Farley and Lucas Sabalka. Presentations of graph braid groups, 2009.

[9] Robert Ghrist. Configuration spaces and braid groups on graphs in robotics. In *Knots, braids, and mapping class groups—papers dedicated to Joan S. Birman (New York, 1998)*, volume 24 of *AMS/IP Stud. Adv. Math.*, pages 29–40. Amer. Math. Soc., Providence, RI, 2001.

[10] L. Sabalka. *Braid groups on graphs*. PhD thesis, University of Illinois at Urbana-Champaign, 2006.